Using Defeasible Information to Obtain Coherence*

Giovanni Casini† Thomas Meyer‡

March 26, 2017

Abstract

We consider the problem of obtaining coherence in a propositional knowledge base using techniques from Belief Change. Our motivation comes from the field of formal ontologies where coherence is interpreted to mean that a concept name has to be satisfiable. In the propositional case we consider here, this translates to a propositional formula being satisfiable. We define belief change operators in a framework of nonmonotonic preferential reasoning. We show how the introduction of defeasible information using contraction operators can be an effective means for obtaining coherence.

1 Introduction

Consider a scenario in which we are presented with the following information about red blood cells: *Avian red blood cells and mammalian red blood cells are vertebrate red blood cells, and vertebrate red blood cells have a cell membrane. Also, vertebrate red blood cells have a nucleus, but mammalian red blood cells don't.* In a propositional knowledge base, this information can be represented as follows: $\{a \rightarrow v, m \rightarrow v, v \rightarrow n, m \rightarrow \neg n\}$. An obvious problem with this formulation is the consequence that mammalian red blood cells don't exist ($m \rightarrow \bot$). We can be reasonably certain that this type of modelling is not what was intended. In the Description Logic (DL) community the analogous problem (i.e. when DL concept names are not satisfiable) is referred to as *incoherence* [15], and we shall adopt the same terminology for the propositional context. Our goal here is to consider a generalised version of coherence for which a designated complex concept has to be satisfiable. In the propositional case this translates to ensuring that a designated formula has to be satisfiable. The techniques we develop can be extended to deal with the coherence of a knowledge base (all atoms have to be satisfiable), or requiring that a selected set of formulas be simultaneously satisfiable w.r.t. a

*Pre-print version of the paper has been published by **AAAI Press** in the volume **Proceedings of the 15th International Conference on Principle of Knowledge Representation and Reasoning (KR-16)**. The published version can be purchased and downloaded at: http://www.aaai.org/Press/Proceedings/kr16.php

†Université du Luxembourg, Luxembourg. email: giovanni.casini@uni.lu

‡CAIR (CSIR and University of Cape Town), South Africa. email: tmeyer@cs.uct.ac.za

knowledge base.
The standard approach for achieving coherence is closely related to classical belief base contraction [8]: coherence w.r.t. the atom p is obtained via a base contraction of the formula $\neg p$, thereby ensuring the satisfiability of p. In this paper we investigate mechanisms for obtaining propositional coherence by applying techniques developed in the area of Belief Change, but unlike the standard approaches our version of belief change is applied to a logic for defeasible reasoning. The motivation for our approach is based on the observation that the reason for incoherence in a knowledge base can frequently be traced back to a lack of expressivity in the logic used for representing the knowledge. For example, when stating that vertebrate red blood cells have a nucleus, what the knowledge engineers probably had in mind is that vertebrate red cells *usually* have a nucleus, with mammalian red blood cells being an exception. While (classical) propositional logic is not able to deal with exceptions, there are a range of nonmonotonic formalisms capable of doing so, thereby enabling us to obtain coherence. We consider the approach, mainly developed by Lehmann and Magidor [13], in which conditionals of the form $C \mid\sim D$ (*defeasible information*) are added to the language of propositional logic (with C and D being classical propositional formulas). In this enriched language, equipped with the entailment relation known as Rational Closure, the information about vertebrate blood cells usually having a nucleus would be expressible as $v \mid\sim n$. This shift allows us to restore coherence by substituting selected bits of classical information with their defeasible counterparts instead of eliminating them completely, as most approaches do. In our example above, substituting $v \rightarrow n$ (vertebrate red blood cells have a nucleus) with $v \mid\sim n$ (vertebrate red blood cells usually have a nucleus) results in a coherent defeasible knowledge base.

2 Background

We consider a finitely generated propositional language $\mathcal{L}$ with lower case letters $a, b, c, \ldots$ denoting atoms, and capital letters $A, B, C, \ldots$ denoting elements of $\mathcal{L}$. We consider the language $\mathcal{L}^{\vdash}$ of (classical) conditionals of the form $C \vdash D$ (for $C, D \in \mathcal{L}$), with a conditional $C \vdash D$ being semantically equivalent to the implication $C \rightarrow D$. It is thus easily seen that the logic based on $\mathcal{L}^{\vdash}$ is exactly as expressive as the propositional logic based on $\mathcal{L}$. We extend $\mathcal{L}^{\vdash}$ by introducing the language $\mathcal{L}^{\mid\sim}$ consisting of all elements of $\mathcal{L}^{\vdash}$, as well as defeasible conditionals of the form $C \mid\sim D$ (for $C, D \in \mathcal{L}$). Classical conditionals can be viewed as propositional versions of classical subsumptions occurring in DL TBoxes [2], while defeasible conditionals are propositional versions of defeasible subsumptions [6]. In this paper we skip the semantic characterisation of defeasible conditionals [13], except to note that classical conditionals can be represented as defeasible ones: $R \Vdash C \vdash D$ iff $R \Vdash C \wedge \neg D \mid\sim \bot$. We sometimes abuse notation by using $C \wedge \neg D \mid\sim \bot$ to refer to classical conditionals. We consider conditional knowledge bases K consisting of finite sets of classical and defeasible conditionals. The notion of entailment that we associate with this semantics is *Rational Closure* (RC) [13]. Let $\overrightarrow{\mathsf{K}} = \{C \rightarrow D \mid C \vdash D \in \mathsf{K} \text{ or } C \mid\sim D \in \mathsf{K}\}$; through a series of classical propositional decision steps over the set $\overrightarrow{\mathsf{K}}$ and the set $\mathcal{A}_{\mathsf{K}}$ of the antecedents of the conditionals in K ($\mathcal{A}_{\mathsf{K}} = \{C \mid C \sqsubsetsim D \in \mathsf{K}\}$) we assign a rank $r_{\mathsf{K}}(C)$

to every formula C, where $r_{\mathsf{K}}(C)$ is either a natural number or ∞. The rank of $C \mid\sim D$ w.r.t. K, $r_{\mathsf{K}}(C \mid\sim D)$, is equal to the rank associated with its antecedent $r_{\mathsf{K}}(C)$. Intuitively, the rank of a defeasible conditional indicates its level of defeasibility: the lower the rank, the more likely we are to discard it. The rank of $C \vdash D$ is ∞, which is also the rank of its defeasible counterpart $C \wedge \neg D \mid\sim \bot$. For $i \in \{0, \ldots, n-1, \infty\}$ we let $\mathsf{K}^i = \{C \mid\sim D \in \mathsf{K} \mid r_{\mathsf{K}}(C \mid\sim D) = i\}$, and we let $\mathsf{K}^{<\infty} = \mathsf{K} \backslash \mathsf{K}^{\infty}$. The RC of K, $\models_{rc}$, is defined as follows: $\mathsf{K} \models_{rc} C \mid\sim D$ iff $r_{\mathsf{K}}(C) < r_{\mathsf{K}}(C \wedge \neg D)$ or $r_{\mathsf{K}}(C) = \infty$. We let $\mathcal{RC}(\mathsf{K}) = \{C \mid\sim D \mid \mathsf{K} \models_{rc} C \mid\sim D\}$. We extend the notion of coherence to conditional knowledge bases: C is coherent for K iff $\mathsf{K} \not\models_{rc} C \vdash \bot$ iff $\mathsf{K} \not\models_{rc} C \mid\sim \bot$. K is *inconsistent* iff $\mathsf{K} \models_{rc} \top \vdash \bot$. Two knowledge bases are *rank equivalent* iff their Rational Closures are equal.

We chose to use RC in this paper for a number of reasons. Most importantly, below we are going to investigate a form of belief contraction of conditionals of the form $C \mid\sim \bot$. It turns out that all the prominent entailment relations proposed in the preferential framework, e.g. [11, 13, 12, 4, 3], are equivalent for such conditionals ($C \mid\sim \bot$ is a consequence of K in one of them iff it is a consequence in all of them). So, if we model a contraction operator for RC, we can immediately apply the same operator to all the other entailment relations in this family. Note that only one half of the Deduction Theorem holds for defeasible conditionals in the context of RC, but the other direction doesn't.

Proposition 1 *If* $\mathsf{K} \models_{rc} C \wedge D \mid\sim E$ *then* $\mathsf{K} \models_{rc} C \mid\sim \neg D \vee E$. *However, it may be that* $\mathsf{K} \not\models_{rc} C \wedge D \mid\sim E$ *when* $\mathsf{K} \models_{rc} C \mid\sim \neg D \vee E$.

So, $C \mid\sim \neg D \vee E$ can be viewed as a weakening of $C \wedge D \mid\sim E$ in the context of RC, a result which forms the crux of the notion of weakening that we will be employing.

3 Simple Weakening

In this section we present an initial proposal for obtaining coherence for a formula C w.r.t. a classical conditional knowledge base $\mathsf{K}^{\vdash}$. In developing the proposal we adopt the techniques used in *kernel-contraction* [8]. A *kernel* for $C \vdash \bot$ in $\mathsf{K}^{\vdash}$ is a set of conditionals $\mathsf{K}' \subseteq \mathsf{K}^{\vdash}$ s.t. $\mathsf{K}' \models C \vdash \bot$ and for every $\mathsf{K}'' \subset \mathsf{K}', \mathsf{K}'' \not\models C \vdash \bot$. We denote the set of kernels for $C \vdash \bot$ in $\mathsf{K}^{\vdash}$ by $\mathsf{K}^{\vdash} \perp\!\!\!\perp (C \mid\sim \bot)$.
To explain the principle behind the proposal, we return to our red blood cell example. Let $\mathsf{K}^{\vdash} = \{a \vdash v, m \vdash v, v \vdash n, m \vdash \neg n\}$, and assume we attempt to obtain coherence for m. To do so we identify the kernels for m. In this case there is one: $\{m \vdash v, v \vdash n, m \vdash \neg n\}$. Classical kernel-contraction would remove some non-empty subset of this kernel, making the resulting set coherent for m. We propose instead to identify those elements of the kernel best suited for weakening and to replace one of them with its defeasible counterpart. To decide which elements are best suited for weakening, we consider the *complete weakening* of $\mathsf{K}^{\vdash}$: $\mathsf{K}^{\mid\sim} = \{C \mid\sim D \mid C \vdash D \in \mathsf{K}^{\vdash}\}$, in our case $\mathsf{K}^{\mid\sim} = \{a \mid\sim v, m \mid\sim v, v \mid\sim n, m \mid\sim \neg n\}$. We use the rank of the elements of the complete weakening to identify those best suited for weakening—the lower the rank, the more suited a conditional is for weakening. In our example the conditional $v \mid\sim n$ has the lowest rank, 0, and we end up with the conditional knowledge base $\mathsf{K}^{-}_{m\vdash\bot} = \{a \vdash v, m \vdash v, v \mid\sim n, m \vdash \neg n\}$. Observe that m is coherent for $\mathsf{K}^{-}_{m\vdash\bot}$.

More generally, consider a classical conditional KB $\mathsf{K}^{\vdash}$ and a formula C. For $\mathsf{K}' \subseteq \mathsf{K}^{\vdash}$, we let $\min_r(\mathsf{K}') = \{D \vdash E \subseteq \mathsf{K}' \mid r_{\mathsf{K}^{\mid\sim}}(D \mid\sim E) \leqslant r_{\mathsf{K}^{\mid\sim}}(F \mid\sim G)$ for every $F \vdash G \in \mathsf{K}'\}$, and we let σ be a choice function $\sigma : \mathscr{P}(K^{\vdash}) \longrightarrow \mathsf{K}^{\vdash}$ such that $\emptyset \subset \sigma(\mathsf{K}') \subseteq \mathsf{K}'$. Furthermore, let $\mathsf{K}^w = \{D \vdash E \mid D \vdash E \in \sigma(\min_r(J))$ for some $J \in \mathsf{K}^{\vdash}\perp\!\!\!\perp(C \mid\sim \bot)\}$. Then we define $\mathsf{K}^{\vdash^{-\sigma}}_{C\vdash\bot}$ as $\mathsf{K}^{\vdash^{-\sigma}}_{C\vdash\bot} := \mathsf{K}^{\vdash}\backslash\mathsf{K}^w \cup \{D \mid\sim E \mid D \vdash E \in \mathsf{K}^w\}$.

Proposition 2 $\mathsf{K}^{\vdash^{-\sigma}}_{C\vdash\bot} \not\models_{rc} C \vdash \bot$ *if and only if for every set* K' *in* $\mathsf{K}^{\vdash}\perp\!\!\!\perp(C \mid\sim \bot)$ *the rank of the elements of* $\min_r(\mathsf{K}')$ *w.r.t.* $\mathsf{K}^{\mid\sim}$ *is less than* ∞.

The procedure for Simple Weakening we describe above is simple, elegant, and easy to implement, since reasoning can be reduced to classical propositional reasoning. Also, this proposal is immediately applicable and implementable for DLs as well. However, it is strongly syntax-dependent, and it does not always obtain coherence. Consider the knowledge base $\{a \wedge \neg v \vdash \bot, m \wedge \neg v \vdash \bot, v \wedge \neg n \vdash \bot, m \wedge n \vdash \bot\}$, in which every formula is logically equivalent to one in our original knowledge base. In this case the defeasible versions are not weaker than their strict counterparts (they all have rank ∞). Hence none of the potential solutions provided by Simple Weakening renders m coherent. In the next section we consider a more nuanced version of weakening which eliminates these drawbacks.

4 Nuanced Weakening

In this section we present Nuanced Weakening as a form of *contraction* of a conditional knowledge base K for which conditionals of the form $C \mid\sim \bot$ are contracted. Unlike Simple Weakening, it is always successful in obtaining coherence, and is more faithful to the Principle of Minimal Change. Like Simple Weakening, it is easily implementable since it can be reduced to classical propositional reasoning.
It can be viewed as a contraction operator on theories (logically closed sets of formulas) since we obtain a weakening of K. That is, we obtain a new knowledge base K' (which is not necessarily a subset of K) s.t. $\mathsf{K} \models_{rc} C \mid\sim D$ for every conditional $C \mid\sim D \in \mathsf{K}'$. But it can also be viewed as a belief base contraction operator once we have applied a form of completion to K. So Nuanced Weakening is, in some sense, a hybrid between theory and base contraction. Our approach has a number of advantages. The fact that information is weakened instead of eliminated ensures that the Principle of Minimal Change is taken into account. Also, a foundationalist approach, has two important advantages. It ensures that the Principle of *Categorical Matching* is satisfied (we start with a base and end with one), making iteration immediately possible, and it eases the move to the application of our techniques for simultaneously obtaining coherence for all concept names in formal ontologies represented as Description Logics. In that context, the representation of ontologies is foundationalist. In fact, optimised software tools for identifying kernels (referred to as *justifications* in the DL community) already exist [9].
First, note that the entailment of statements of the form $C \mid\sim \bot$ depends solely on those elements of K with rank ∞, and can be reduced to classical propositional reasoning.

Proposition 3 $\mathsf{K} \models_{rc} C \mid\sim \bot$ *iff* $\mathsf{K}^{\infty} \models_{rc} C \mid\sim \bot$ *iff* $\overrightarrow{\mathsf{K}^{\infty}} \models \neg C$.

Next we define a form of completion of K that will allow us to use a propositional contraction operator over $\overrightarrow{\mathsf{K}^{\infty}}$ and minimise the loss of information, just implementing a minimal amount of weakening of the information. So, we start with a finite set K of (classical and defeasible) conditionals, and apply the following transformation:

1. We translate all the classical conditionals $C \vdash D$ in K to their defeasible versions $C \wedge \neg D \mid\sim \bot$

2.Then we rewrite the antecedents of all the conditionals in K to be in conjunctive normal form, restricted to the atoms occurring in the antecedent. Every conditional will have the form $\bigwedge \Gamma \mid\sim D$, where Γ is a set of propositional clauses.

3. For each conditional $\bigwedge \Gamma \mid\sim D \in \mathsf{K}$, and for every $\Gamma' \subset \Gamma$ we also add to K the conditional $\bigwedge \Gamma' \mid\sim D \vee \bigvee\{\neg C \mid C \in \Gamma \backslash \Gamma'\}$, where $\bigwedge \varnothing$ is defined as $\top$.

The basic idea behind the above transformation is to add to K all the possible weakenings of a conditional that can be obtained via the application of Proposition 1. For example, $p \vdash q$ is rewritten into $p \wedge \neg q \mid\sim \bot$, and then add to K the weaker defeasible conditionals $p \mid\sim \bot \vee \neg\neg q$, $\neg q \mid\sim \bot \vee \neg p$, and $\top \mid\sim \bot \vee \neg p \vee \neg\neg q$. In the rest of the section we will consider that a KB K is completed as described.

Example 1 *Consider again the knowledge base in our red blood cell example. First, we rewrite* K *as* $\{a \wedge \neg v \mid\sim \bot, m \wedge \neg v \mid\sim \bot, v \wedge \neg n \mid\sim \bot, m \wedge n \mid\sim \bot\}$ *(the antecedents are already in conjunctive normal form). Then we complete the knowledge base by adding the weakened versions of the conditionals to obtain* K'*. For example, for* $m \wedge \neg v \mid\sim \bot$*, we will also add* $m \mid\sim \bot \vee \neg\neg v$*,* $\neg v \mid\sim \bot \vee \neg m$*, and* $\top \mid\sim \bot \vee \neg m \vee \neg\neg v$*. And similarly for the other conditionals in* K*. Observe that these weakened conditionals are equivalent (w.r.t. RC) to* $m \mid\sim v$*,* $\neg v \mid\sim \neg m$*, and* $\top \mid\sim \neg m \vee v$*. To ease readability we perform similar logically equivalent transformations in our examples.*

Later we'll see that only the first two steps of the above completion of K actually need to be performed, and that the third step can be simulated without explicitly introducing any new conditionals into K. We are now ready to define our contraction operator. Assume we have a conditional knowledge base K that has been completed as described above, and a conditional $C \mid\sim \bot$ with which we want to contract. We are going to define a standard *kernel*-base operator with a minor twist. From Proposition 3 we know that we need to define kernels only with respect to K^{∞}. So, let $\mathsf{K}^{\infty} \perp\!\!\!\perp (C \mid\sim \bot)$ be the set of the kernels of $C \mid\sim \bot$ w.r.t. K^{∞}. Following the standard kernel approach, we define an *incision function* σ as a choice function that chooses at least one element in every set in $\mathsf{K}^{\infty} \perp\!\!\!\perp (C \mid\sim \bot)$. In particular, σ must satisfy the following conditions to be an incision function ([8], Definition 2.30):

1. $\sigma(\mathsf{K}^{\infty} \perp\!\!\!\perp (C \mid\sim \bot)) \subseteq \bigcup[(\mathsf{K}^{\infty} \perp\!\!\!\perp (C \mid\sim \bot))$;
2. If $\varnothing \neq X \in \mathsf{K}^{\infty} \perp\!\!\!\perp (C \mid\sim \bot)$, then $X \cap \sigma(\mathsf{K}^{\infty} \perp\!\!\!\perp (C \mid\sim \bot)) \neq \varnothing$.

We let the Nuanced Weakening of K for C be the result obtained by the contraction operator $\barwedge$ which is defined as $\mathsf{K}^{\barwedge}_{(C \mid\sim \bot)} = \mathsf{K} \backslash \sigma(\mathsf{K}^{\infty} \perp\!\!\!\perp (C \mid\sim \bot))$.

Example 2 *Consider the knowledge* K' *from Example 1. The conditionals below each line are weaker than the ones above, and their ranks are indicated with a superscript:*

$$\dfrac{\dfrac{a \wedge \neg v \mathrel{|\!\sim} \bot^{\infty}}{a \mathrel{|\!\sim} v^{0} \quad \neg v \mathrel{|\!\sim} \neg a^{0}}}{\top \mathrel{|\!\sim} \neg a \vee v^{0}} \qquad \dfrac{\dfrac{\boxed{m \wedge \neg v \mathrel{|\!\sim} \bot^{\infty}}}{\boxed{m \mathrel{|\!\sim} v^{\infty}} \quad \neg v \mathrel{|\!\sim} \neg m^{0}}}{\top \mathrel{|\!\sim} \neg m \vee v^{0}}$$

$$\dfrac{\dfrac{\boxed{v \wedge \neg n \mathrel{|\!\sim} \bot^{\infty}}}{v \mathrel{|\!\sim} n^{0} \quad \neg n \mathrel{|\!\sim} \neg v^{0}}}{\top \mathrel{|\!\sim} \neg v \vee n^{0}} \qquad \dfrac{\dfrac{\boxed{m \wedge n \mathrel{|\!\sim} \bot^{\infty}}}{\boxed{m \mathrel{|\!\sim} \neg n^{\infty}} \quad n \mathrel{|\!\sim} \neg m^{0}}}{\top \mathrel{|\!\sim} \neg m \vee \neg n^{0}}$$

The boxed conditionals are the ones ocurring in some kernel for $m \mathrel{|\!\sim} \bot$. There are four possible kernels, each containing three conditionals: $v \wedge \neg n \mathrel{|\!\sim} \bot$ and exactly one conditional from each of the remaining two groups. An incision function σ will choose at least one element from each kernel, eliminating it from the knowledge base.

In propositional logic kernel contraction is characterised by a set of postulates: *Success, Inclusion, Core-retainment,* and *Uniformity* ([8], Theorem 2.32). These postulates can be reformulated for our contraction operation $\barwedge$ as follows:

• *Success:* If $C \mathrel{|\!\sim} \bot \notin \mathcal{RC}(\varnothing)$ then $C \mathrel{|\!\sim} \bot \notin \mathcal{RC}(\mathsf{K}^{\barwedge}_{C\mathrel{|\!\sim}\bot})$;

• *Inclusion:* $\mathsf{K}^{\barwedge}_{C\mathrel{|\!\sim}\bot} \subseteq \mathsf{K}$;

• *Core-retainement:* If $D \mathrel{|\!\sim} E \in \mathsf{K}$ and $D \mathrel{|\!\sim} E \notin \mathsf{K}^{\barwedge}_{C\mathrel{|\!\sim}\bot}$, then there is a K' s.t. $\mathsf{K}' \subseteq \mathsf{K}$ and $C \mathrel{|\!\sim} \bot \notin \mathcal{RC}(\mathsf{K}')$ and $C \mathrel{|\!\sim} \bot \in \mathcal{RC}(\mathsf{K}' \cup \{D \mathrel{|\!\sim} E\})$;

• *Uniformity:* If for all $\mathsf{K}' \subseteq \mathsf{K}$ it holds that $C \mathrel{|\!\sim} \bot \in \mathcal{RC}(\mathsf{K}')$ iff $D \mathrel{|\!\sim} \bot \in \mathcal{RC}(\mathsf{K}')$, then $\mathsf{K}^{\barwedge}_{C\mathrel{|\!\sim}\bot} = \mathsf{K}^{\barwedge}_{D\mathrel{|\!\sim}\bot}$;

where K is a conditional KB completed as described above. It is easy to prove the correspondent representation theorem for our conditional contraction.

Proposition 4 (Representation Theorem for Kernel Conditional Contraction.) *The operator $\barwedge$ is an of operator that satisfies* Success, Inclusion, Core-retainement, *and* Uniformity*, and every operator satisfying such properties can be modelled as a $\barwedge$-operator.*

Nuanced Weakening overcomes some of the limits that we have pointed out about Simple Weakening, but it has some problems as well. The completion procedure is quite cumbersome, and the kind of contraction we have presented, a simple kernel-contraction, despite being easily implementable (especially in view of an adaptation for Description Logics), is not really satisfying from the point of view of the Principle of Minimal Change. There are various refinement that we can implement in order to overcome such limits, that we will not be able to properly present here.
Dealing with the issue about the complexity of the completion procedure, we can actually ease it a lot, avoiding the actual construction and introduction in the KB of new conditionals. As mentioned in the *Background* section, we can decide the ranking of every conditional in K using $\overrightarrow{\mathsf{K}}$ and the set $\mathcal{A}_{\mathsf{K}} = \{\bigwedge\Gamma \mid \bigwedge\Gamma \precsim D \in \mathsf{K}\}$, containing the antecedents of the conditionals in K. Once that every formula in $\mathcal{A}_{\mathsf{K}}$ is into its conjunctive normal form $\bigwedge\Gamma$, we can prove that we can actually 'simulate' the kernel contraction described above by extending $\mathcal{A}_{\mathsf{K}}$ with the formulas $\bigwedge\Theta$, with $\Theta \subset \Gamma$ for some $\bigwedge\Gamma \in \mathcal{A}_{\mathsf{K}}$. Once we do such a step, it is also possible to refine this kernel-contraction procedure in order minimising the loss of information, up to defining a

maxichoice contraction operator [1], that from the point of view of minimising change is the most appealing approach. When dealing with theories (closed under entailment) such an approach can give back counterintuitive results [1], but when working with finite bases this approach is considered desirable [8].
Finally, the present contraction operations over conditionals $C \precsim \bot$, beyond that for dealing with the problem of coherence, can be used to define a general *revision operator* over a conditional KB, that is, a class of operators that, given a conditional KB K and a conditional $C \precsim D$, allows to define a new KB $\mathsf{K}^*_{C \precsim D}$ that guarantees to contains $C \precsim D$ while preserving logical consistency.

5 Conclusion

In this paper we combine formal tools taken from the Belief Revision and the Non-monotonic Reasoning communities in order to define belief change operators that are appropriate for enforcing coherence in conditional knowledge bases, minimising the loss of information. The proposed operators are easily implementable. Even though the procedures presented here are not computationally tractable, from a practical perspective Horridge [9] has shown that computing kernels is frequently feasible even for large DL ontologies.

With the exception of the work of Kern-Isberner [10], which is not directly relevant to ours, we are not aware of any work about the dynamics of conditional knowledge.

Our proposal can be used to develop a class of belief revision operators for conditional knowledge bases on the basis of the introduced contraction operators. Also, we plan to extend these results to Description Logics (DLs). Much of the content of this paper is already applicable to DLs. Furthermore, the decision procedures for the main inference operations in the framework of preferential reasoning have already been defined for DLs [14, 6, 4, 3, 7, 5]. So the only remaining obstacle for a DL version of Nuanced Weakening is the definition of a proper normal form to be used in substitution of the version of conjunctive normal form used in Nuanced Weakening. Beyond that, a proper analysis of iterated and multiple base contraction is needed, and an analysis of the same belief change problems in the framework of logically closed theories, more in line with an AGM-style analysis is also possible.

Acknowledgments. The work of Giovanni Casini has been supported by the Fonds National de la Recherche, Luxembourg, and cofunded by the Marie Curie Actions of the European Commission (FP7-COFUND) (AFR/9181001).

References

[1] C. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.

[2] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider, editors. *The Description Logic Handbook*. Cambridge University Press, 2 edition, 2007.

[3] G. Casini, T. Meyer, K. Moodley, and R. Nortje. Relevant closure: A new form of defeasible reasoning for description logics. In *Proc. of JELIA 2014*, pages 92–106. Springer, 2014.

[4] G. Casini and U. Straccia. Defeasible inheritance-based description logics. *JAIR*, 48:415–473, 2013.

[5] G. Casini and U. Straccia. Lexicographic closure for defeasible description logics. forthcoming.

[6] L. Giordano, N. Olivetti, V. Gliozzi, and G.L. Pozzato. A non-monotonic description logic for reasoning about typicality. *Artif. Intell.*, 195:165–202, 2013.

[7] L. Giordano, N. Olivetti, V. Gliozzi, and G.L. Pozzato. Semantic characterization of rational closure: From propositional logic to description logics. *Artif. Intell.*, 226:1–33, 2015.

[8] S.O. Hansson. *A Textbook of Belief Dynamics: Theory Change and Database Updating*. Kluwer, 1999.

[9] M. Horridge. *Justification based explanation in ontologies*. The University of Manchester, 2011.

[10] Gabriele Kern-Isberner. Linking iterated belief change operations to nonmonotonic reasoning. In Gerhard Brewka and Jérôme Lang, editors, *Proceedings of KR 2008*, pages 166–176. AAAI Press, 2008.

[11] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artif. Intell.*, 44:167–207, 1990.

[12] D. Lehmann. Another perspective on default reasoning. *Ann. of Math. and Artif. Intell.*, 15(1):61–82, 1995.

[13] D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artif. Intell.*, 55:1–60, 1992.

[14] T. Łukasiewicz. Expressive probabilistic description logics. *Artif. Intell.*, 172:852 – 883, 2008.

[15] Guilin Qi and Anthony Hunter. Measuring incoherence in description logic-based ontologies. In *The Semantic Web*, volume 4825 of *LNCS*, pages 381–394. Springer, 2007.